

Lifting from the Deep: Convolutional 3D Pose Estimation from a Single Image

Denis Tome
University College London
D.Tome@cs.ucl.ac.uk

Chris Russell
The Turing Institute
crussell@turing.ac.uk

Lourdes Agapito
University College London
l.agapito@cs.ucl.ac.uk

<http://visual.cs.ucl.ac.uk/pubs/liftingFromTheDeep>

Abstract

We propose a unified formulation for the problem of 3D human pose estimation from a single raw RGB image that reasons jointly about 2D joint estimation and 3D pose reconstruction to improve both tasks. We take an integrated approach that fuses probabilistic knowledge of 3D human pose with a multi-stage CNN architecture and uses the knowledge of plausible 3D landmark locations to refine the search for better 2D locations. The entire process is trained end-to-end, is extremely efficient and obtains state-of-the-art results on Human3.6M outperforming previous approaches both on 2D and 3D errors.

1. Introduction

Estimating the full 3D pose of a human from a single RGB image is one of the most challenging problems in computer vision. It involves tackling two inherently ambiguous tasks. First, the 2D location of the human joints, or landmarks, must be found in the image, a problem plagued with ambiguities due to the large variations in visual appearance caused by different camera viewpoints, external and self occlusions or changes in clothing, body shape or illumination. Next, lifting the coordinates of the 2D landmarks into 3D from a single image is still an ill-posed problem – the space of possible 3D poses consistent with the 2D landmark locations of a human, is infinite. Finding the correct 3D pose that matches the image requires injecting additional information usually in the form of 3D geometric pose priors and temporal or structural constraints.

We propose a new joint approach to 2D landmark detection and full 3D pose estimation from a single RGB image that takes advantage of reasoning jointly about the estimation of 2D and 3D landmark locations to improve both tasks. We propose a novel CNN architecture that learns to combine the image appearance based predictions provided by *convolutional-pose-machine* style 2D landmark detectors [43], with the geometric 3D skeletal information en-

coded in our novel pre-trained model of 3D human pose. Our manifold is learned exclusively from 3D mocap

The information captured by our 3D human pose model is embedded in the CNN architecture as an additional layer that lifts 2D landmark coordinates into 3D while imposing that they lie on the space of physically plausible poses. The advantage of integrating the output proposed by the 2D landmark location predictors – based purely on image appearance – with the 3D pose predicted by a probabilistic model, is that the 2D landmark location estimates are improved by guaranteeing that they satisfy the anatomical 3D constraints encapsulated in the human 3D pose model. In this way, both tasks clearly benefit from each other.

A further advantage of our approach is that our 2D and 3D training data sources may be completely independent. Our deep architecture only needs that images are annotated with 2D poses, not 3D poses. The human pose model is trained independently and exclusively from 3D mocap data. This decoupling between 2D and 3D training data presents a huge advantage since we can augment our training sets completely independently. For instance we can take advantage of extra 2D pose annotations without the need for 3D ground truth or extend the 3D training data to further mocap datasets without the need for synchronized 2D images.

Our contribution: In this work, we show how to integrate a pre-learned 3D human pose model directly within a novel CNN architecture (illustrated in figure 1) for joint 2D landmark and 3D human pose estimation. In contrast to pre-existing methods, we do not take a pipeline approach that takes 2D landmarks as given. Instead, we show how such a model can be used as part of the CNN architecture itself, and how the architecture can learn to use physically plausible 3D reconstructions in its search for better 2D landmark locations. Our method achieves state-of-the-art results on the Human3.6M dataset both in terms of 2D and 3D errors.

2. Related Work

We first describe methods that assume that 2D joint locations are provided as input and focus on solving the 3D

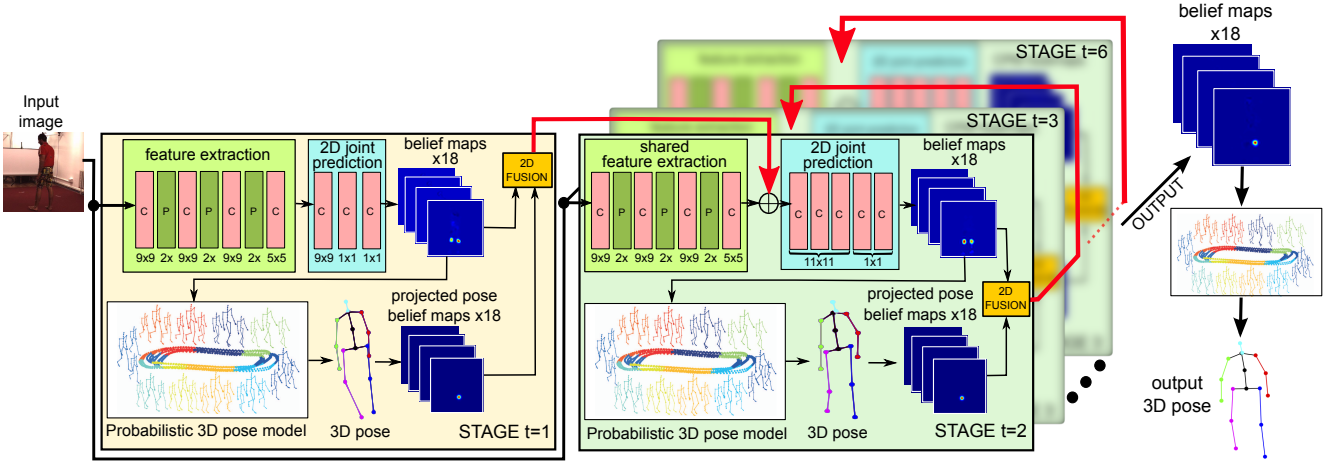


Figure 1: Our multistage deep architecture for 2D/3D human pose estimation. Each stage t produces as output a set of belief maps for the location of the 2D landmarks (one per landmark). The belief maps from each stage t , combined with learned image features, are used as input to the next stage $t + 1$ (red arrows show information flowing from one stage to the next). Internally, each stage learns to combine: (a) appearance-based belief maps provided by convolutional 2D landmark predictors, with (b) projected pose belief maps, proposed by our new probabilistic 3D pose model that encodes 3D structural information. The 3D pose layer is responsible for lifting 2D landmark coordinates into 3D and projecting them onto the space of valid 3D poses. These two belief maps are then fused into a single set of output proposals for the 2D landmark locations per stage. The accuracy of the 2D and 3D landmark locations increases progressively through the stages. The overall architecture is trained end-to-end using back-propagation. [Best viewed in color.]

lifting problem and follow with methods that learn to estimate the 3D pose directly from images.

3D pose from known 2D joint positions: A large body of work has focused on recovering the 3D pose of people given perfect 2D joint positions as input. Early approaches [18, 33, 24, 6] took advantage of anatomical knowledge of the human skeleton or joint angle limits to recover pose from a single image. More recent methods [13, 27, 3] have focused on learning a prior statistical model of the human body directly from 3D mocap data.

Non-rigid structure from motion approaches (NRSfM) also recover 3D articulated motion [8, 4, 14, 19] given known 2D correspondences for the joints in every frame of a monocular video. Their huge advantage, as unsupervised methods, is they do not need 3D training data, instead learning a linear basis for the 3D poses purely from 2D data. Their main drawback is their need for significant camera movement throughout the sequence to guarantee accurate 3D reconstruction. Recent work on NRSfM applied to human pose estimation has focused on escaping these limitations by the use of a linear model to represent shape variations of the human body. For instance, [10] defined a generative model based on the assumption that complex shape variations can be decomposed into a mixture of primitive shape variations and achieve competitive results.

Representing human 3D pose as a linear combination of a sparse set of 3D bases, pre-trained using 3D mocap data, has also proved a popular approach for articulated human motion [27, 42, 47]. While [47] propose a convex relaxation to jointly estimate the coefficients of the sparse representation and the camera viewpoint [27] and [42] enforce limb

length constraints. Although these approaches can reconstruct 3D pose from a single image, their best results come from imposing temporal smoothness on the reconstructions of a video sequence.

Recently, Zhao *et al.* [45] achieved state-of-the-art results by training a simple neural network to recover 3D pose from known 2D joint positions. While the results on perfect 2D input data are impressive, the inaccuracies in 2D joint estimation are not modeled and the performance of this approach combined with joint detectors is unknown.

3D pose from images: Most approaches to 3D pose inference directly from images fall into one of two categories: (i) models that learn to regress the 3D pose directly from image features and (ii) pipeline approaches where the 2D pose is first estimated, typically using discriminatively trained part models or joint predictors, and then lifted into 3D. While regression based methods suffer from the need to annotate all images with ground truth 3D poses – a technically complex and elaborate process – for pipeline approaches the challenge is how to account for uncertainty in the measurements. Crucial to both types of approaches is the question of how to incorporate the 3D dependencies between the different body joints or to leverage other useful 3D geometric information in the inference process.

Many earlier works on human pose estimation from a single image relied on discriminatively trained models to learn a direct mapping from image features such as silhouettes, HOG or SIFT, to 3D human poses without passing through 2D landmark estimation [1, 12, 11, 23, 31].

More recent direct approaches make use of deep learning [20, 21, 39, 40]. Regression-based approaches train an

end-to-end network to predict 3D joint locations directly from the image [40, 20, 21, 46]. While [21] incorporate model joint dependencies in the CNN via a max-margin formalism, others [46] impose kinematic constraints by embedding a differentiable kinematic model into the deep learning architecture. More recently, Tekin *et al.* [34] propose a deep regression architecture for structured prediction that combines traditional CNNs for supervised learning with an auto-encoder that implicitly encodes 3D dependencies between body parts.

As CNNs have become more prevalent, 2D joint estimation [43] has become increasingly reliable and many recent works have looked to exploit this using a pipeline approach. Papers such as [9, 16, 39, 25] first estimate 2D landmarks and later 3D spatial relationships are imposed between them using structured learning or graphical models.

Simo-Serra *et al.* [32] were one of the first to propose an approach that naturally copes with the noisy detections inherent to off-the-shelf body part detectors by modeling their uncertainty and propagating it through 3D shape space while guaranteeing that geometric and kinematic 3D constraints were satisfied. The work [30] also estimates the location of 2D joints before predicting 3D pose using appearance and the probable 3D pose of discovered parts using a hierarchical Bayesian non-parametric model. Another recent example is Bogo *et al.* [7], who fit a detailed statistical 3D body model [22] to 2D joint proposals.

Zhou *et al.* [48] tackles the problem of 3D pose estimation for a monocular image sequence integrating 2D, 3D and temporal information to account for uncertainties in the model and the measurements. Similar to our proposed approach, Zhou *et al.*'s method [48] does not need synchronized 2D-3D training data, *i.e.* it only needs 2D pose annotations to train the CNN joint regressor and a separate 3D mocap dataset to learn the 3D sparse basis. Unlike our approach, it relies on temporal smoothness for its best performance, and performs poorly on a single image.

Finally, Wu *et al.* [?]’s 3D Interpreter Network, a recent approach to estimate the skeletal structure of common objects (chairs, sofas, ...) bears similarities with our method. While our approaches share common ground in the decoupling of 3D and 2D training data and the use of projection from 3D to improve 2D predictions the network architectures are very different and, unlike us, they do not carry out a quantitative evaluation on the problem of 3D human pose estimation problem.

3. Our Architecture

Figure 1 illustrates the main contribution of our approach, a new multi-stage CNN architecture that can be trained end-to-end to estimate jointly 2D and 3D joint locations. Crucially it includes a novel layer, based on our probabilistic 3D model of human pose, responsible for lift-

ing 2D poses into 3D and propagating 3D information about the skeletal structure to the 2D convolutional layers. In this way, the prediction of 2D pose benefits from the 3D information encoded. Section 4 describes our new probabilistic 3D model of human pose, trained on a dataset of 3D mocap data. Section 5 describes all the new components and layers of our CNN architecture. Finally, Section 6 describes our experimental evaluation on the Human3.6M dataset where we obtain state-of-the-art results. In addition we show qualitative results on images from the MPII dataset.

4. Our Probabilistic 3D Model of Human Pose

One fundamental challenge in creating models of human poses lies in the lack of access to 3D data of sufficient variety to characterize the space of human poses. To compensate for this lack of data we identify and eliminate confounding factors such as rotation in the ground plane, limb length, and left-right symmetry that lead to conceptually similar poses being unrecognized in the training data.

We eliminate some factors by simple pre-processing. Variance due to limb-length is addressed by normalizing the data such that the sum of squared limb lengths on the human skeleton is one; while left-right symmetry is exploited by flipping each pose in the x-axis and re-annotating left as right and vice-versa.

4.1. Aligning 3D Human Poses in the Training Set

Allowing for rotational invariance in the ground-plane is more challenging and requires integration with our data model. We seek the optimal rotations for each pose such that after rotating the poses they are closely approximated by a low-rank compact Gaussian distribution.

We formulate this as a problem of optimization over a set of variables. Given a set of N training 3D poses, each represented as a $(3 \times L)$ matrix \mathbf{P}_i of 3D landmark locations, where $i \in \{1, 2, \dots, N\}$ and L is the number of human joints/landmarks; we seek global estimates of an average 3D pose μ , a set of J orthonormal basis matrices¹ \mathbf{e} and noise variance σ , alongside per sample rotations R_i and basis coefficients a_i such that the following estimate is minimized

$$\arg \min_{\mathbf{R}, \mu, a, \mathbf{e}, \sigma} \sum_{i=1}^N (\|\mathbf{P}_i - \mathbf{R}_i (\mu + a_i \cdot \mathbf{e})\|_2^2 + \sum_{j=1}^J (a_{i,j} \cdot \sigma_j)^2 + \ln \sum_{j=1}^J \sigma_j^2) \quad (1)$$

Where $a_i \cdot \mathbf{e} = \sum_j a_{i,j} \mathbf{e}_j$ is the tensor analog of a multiplication between a vector and a matrix, and $\|\cdot\|_2^2$ is the

¹When we say \mathbf{e} is a set of orthonormal basis matrices we mean that each matrix, if unwrapped into a vector, is of unit norm and orthogonal to all other unwrapped matrices.

squared Frobenius norm of the matrix. Here the y -axis is assumed to point up, and the rotation matrices R_i considered are ground plane rotations of the form

$$R_i = \begin{pmatrix} \cos \theta_i & 0 & \sin \theta_i \\ 0 & 1 & 0 \\ -\sin \theta_i & 0 & \cos \theta_i \end{pmatrix} \quad (2)$$

With the large number of 3D pose samples considered (of the order of 1 million when training on the Human3.6M dataset [15]), and the complex inter-dependencies between samples for \mathbf{e} and σ , the memory requirements mean that it is not possible to solve directly as a joint optimization over all variables using a non-linear solver such as Ceres [2]. Instead, we carefully initialize and then alternate between performing closed-form PPCA [37] to update estimates of μ , \mathbf{a} , \mathbf{e} , σ ; and updating the rotations R_i using Ceres [2] to minimize the above error. As we do this, we steadily increase the size of the basis from 1 through to its target size J . This stops apparent deformations that could be resolved through rotations from becoming locked into the basis at an early stage, and empirically leads to lower cost solutions.

To initialize we use a variant of the Tomasi-Kanade [38] algorithm to estimate the mean 3D pose μ . As the y component is not altered by planar rotations, we take as our estimate of the y component of μ , the mean of each point in the y direction. For the x and z components, we interleave the x and z components of each sample and concatenate them into a large $2N \times L$ matrix \mathbf{M} , and find the rank two approximation of this such that $\mathbf{M} \approx \mathbf{A} \cdot \mathbf{B}$. We then calculate $\hat{\mathbf{A}}$ by replacing each adjacent pair of rows of \mathbf{A} with the closest orthonormal matrix of rank two, and take $\hat{\mathbf{A}}^\dagger \mathbf{M}$ as our estimate² of the x and z components of μ .

The end result of this optimization is a compact low-rank approximation of the data in which all reconstructed poses appear to have the same orientation (see figure 2). In the next section we extend our model to be described as a multi-modal distribution to better capture the variations in the space of 3D human poses.

4.2. A Multi-Modal Model of 3D Human Pose

Although it is possible to directly use the learned Gaussian model in the previous section 4.1 to estimate the 3D (see results in section 6), inspection of figure 2 shows that the data is not-Gaussian distributed and is better described using a multi-modal distribution. In doing this, we are heavily inspired both by approaches such as [26] which characterize the space of human poses as a mixture of PCA bases, and by related works such as [41, 8] that represent poses as an interpolation between exemplars. These approaches are extremely good at modeling tightly distributed poses (e.g. walking) where samples in the testing data are likely to be

² \mathbf{A}^\dagger being the pseudo-inverse of \mathbf{A} .




Figure 2: Visualization of the 3D training data after alignment (see section 4.1) using 2D PCA. Notice how all poses have the same orientation. *Standing-up* poses a), b), c) and d) are all close to each other and far from *sitting-down* poses f) and h) which form another clear cluster.

close to poses seen in training. This is emphatically not the case in much of the Human3.6M dataset, which we use for evaluation. Zooming in on the edges of Figure 2 reveals many isolated paths where motions occur once and then are never revisited again.

Nonetheless, it is precisely these regions of low-density that we are interested in modeling. As such, we seek a coarse representation of the pose space that says something about the regions of low density but also characterizes the multi-modal nature of the pose space. As such, we choose to represent our data as a mixture of probabilistic PCA models using few clusters, and trained using the EM-algorithm [37]. When using a small number of clusters, it is important to initialize the algorithm correctly, as accidentally initializing with multiple clusters about a single mode, can lead to poor density estimates. To initialize our clusters we make use of a simple heuristic.

We first sub-sample the aligned poses (which we refer to as P), and then compute the Euclidean distance d among pairs. We seek a set of k samples S such that the distance between points and their nearest sample is minimized

$$\arg \min_S \sum_{p \in P} \min_{s \in S} d(s, p) \quad (3)$$

We find S using greedy selection, holding our previous estimate of S constant, and iteratively selecting the next candidate s such that $\{s\} \cup S$ minimizes the above cost. A selection of 3D pose samples found using this procedure can be seen in the rendered poses of Figure 2. In practice, we

stop proposing candidates when they occur too close to the existing candidates, as shown by samples (a–d), and only choose one candidate from the dominant mode.

Given these candidates for cluster centers, we assign each aligned point to a cluster representing its nearest candidate and then run the EM algorithm of [37], building a mixture of probabilistic PCA bases.

5. A New Convolutional Architecture for 2D and 3D Pose Inference

Our 3D pose inference from a single RGB image makes use of a multistage deep convolutional architecture, trained end-to-end, that repeatedly fuses and refines 2D and 3D poses, and a second module which takes the final predicted 2D landmarks and lifts them one last time into 3D space for our final estimate (see Figure 1).

At its heart, our architecture is a novel refinement of the Convolutional Pose Machine of Wei *et al.* [43], who reasoned exclusively in 2D, and proposed an architecture that iteratively refined 2D pose estimations of landmarks using a mixture of knowledge of the image and of the estimates of landmark locations of the previous stage. We modify this architecture by generating, at each stage, projected 3D pose belief maps which are fused in a learned manner with the standard maps. From an implementation point of view this is done by introducing two distinct layers, the *probabilistic 3D pose layer* and the *fusion layer* (see figure 1).

Figure 3 shows how the 2D uncertainty in the belief maps is reduced at each stage of the architecture and how the accuracy of the 3D poses increases throughout the stages.

5.1. Architecture of each stage

Our entire architecture consists of 6 stages. Each stage in our sequential architecture is made out of 4 clear components (see figure 1):

Predicting CNN-based belief-maps: we use a set of convolutional and pooling layers, equivalent to those used in the original CPM architecture [43], that combine evidence obtained from image learned features with the belief maps obtained from the previous stage ($t - 1$) to predict an updated set of belief maps for the 2D human joint positions. This step is described in section 5.2.

Lifting 2D belief-maps into 3D: the output of the CNN-based belief maps is then input to a new layer that uses our new pre-trained *probabilistic 3D human pose model* to lift the proposed 2D poses into 3D. The process of lifting 2D into 3D is described in section 5.3.

Projected 2D pose belief maps: The 3D pose estimated by our 3D pose inference layer is then projected back onto the image plane to produce a new set of projected pose belief maps. These belief maps encapsulate 3D dependencies be-

tween the body parts. We describe the projection from 3D to 2D in section 5.4.

2D Fusion layer: The final layer in each stage (described in section 5.5) learns the weights to fuse the two sets of belief maps into a unique set which is then input into the next stage $t + 1$ of the architecture.

Our novel layers were implemented as an extension of the published code of Convolutional Pose Machines [43] inside the *caffe framework* as python layers, and with all weights updated using Stochastic Gradient Descent with momentum.

5.2. Predicting CNN-based belief-maps

Convolutional Pose Machines [43] can be understood as an updating of the earlier work of Ramakrishna *et al.* [28] to use a deep convolutional architecture. In both approaches, at each stage t and for each landmark p , the algorithm returns dense per pixel belief maps $b_t^p[u, v]$, which show how confident it is that a joint center or landmark occurs in any given pixel (u, v) . See Figure 4a for a visualization of stage 1 maps. For stages $t \in \{2, \dots, T\}$ the belief maps are a function of not just the information contained in the image but also the information computed by the previous stage.

In the case of convolutional pose machines, and in our work which uses the same architecture, a summary of the convolution widths and architecture design is shown in Figure 1, with more details of training given in Convolutional Pose Machines (CPM) [43].

Both approaches [43, 28] predict the locations of different landmarks to those captured in the Human3.6M dataset. As such the input and output layers in each stage of the architecture are replaced with a larger set to account for the greater number of joints. The new architecture is then initialized by using the weights with those found in CPM’s model for all pre-existing layers with the new layers randomly initialized.

After retraining, convolutional pose machines return per-pixel estimates of landmark locations, while our techniques for 3D estimation (described in the next section) make use of 2D locations. To transform these belief maps into locations, we select the most confident pixel as the location of each landmark

$$Y_p = \arg \max_{(u,v)} b_p[u, v] \quad (4)$$

5.3. Lifting 2D belief-maps into 3D

We follow [48] in assuming a weak perspective model, and first describe the simplest case of estimating the 3D pose of a single frame using a uni-modal Gaussian 3D pose model as described in section 4. This model is composed of a mean shape μ , a set of basis matrices e and variances

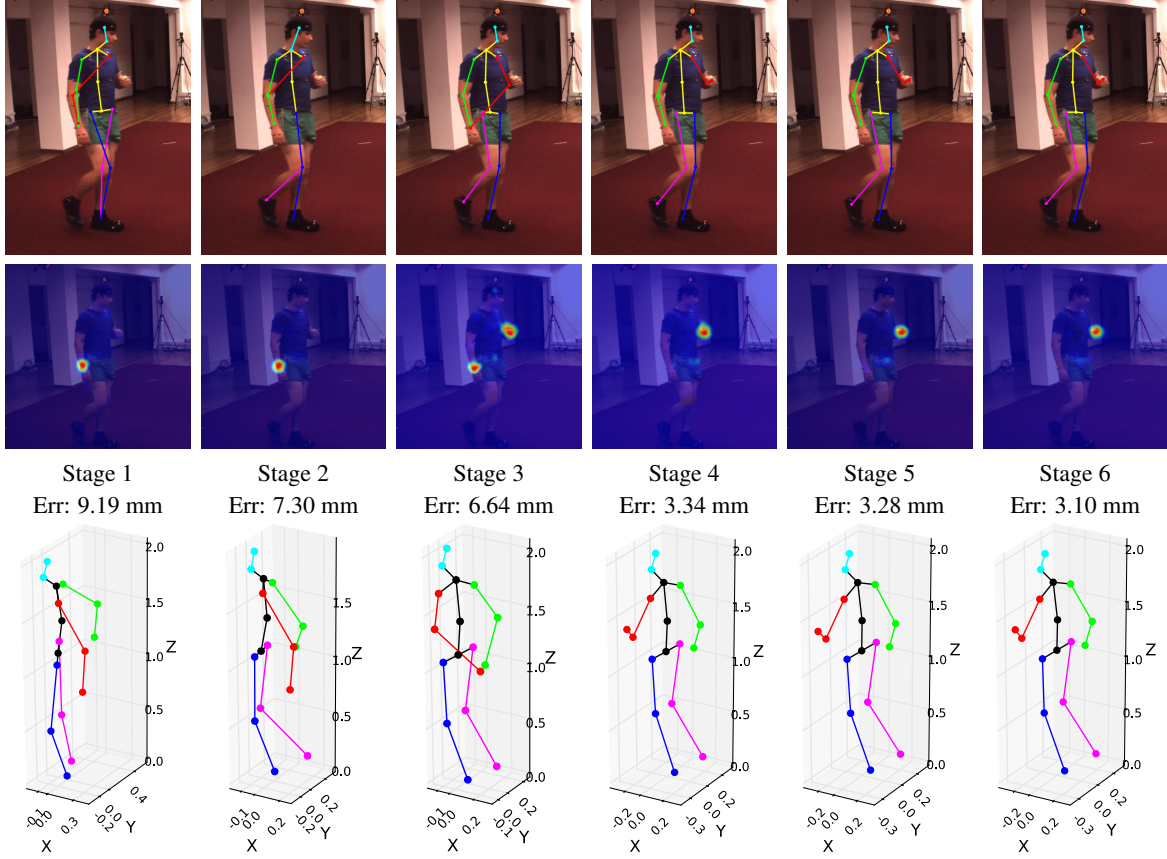


Figure 3: Results returned by different stages of the architecture. *Top*: Evolution of the 2D skeleton after projecting the 3D points back into the 2D space; *Center*: Evolution of the beliefs for the landmark *Left hand* through the stages. *Bottom*: 3D skeleton with the relative mean error per landmark in millimeters. Even when the landmark locations are incorrect, the model returns a physically plausible solution.




Figure 4: Evolution of the beliefs in stage 1 of the architecture for landmark *Spine*. Beliefs are fused together by the *fusion layer*: a) Initial convolutional belief; b) Projected Pose estimation; c) Fused information. See section 5.4.




Figure 5: Landmark refinement: *Left*: 2D predicted landmark positions; *Right*: improved predictions using the projected 3D pose.

σ^2 , and from this we can compute the most probable sample from the model that could give rise to a projected image.

$$\arg \min_{R,a} ||Y - s\Pi ER(\mu + a \cdot \mathbf{e})||_2^2 + ||\sigma \cdot a||_2^2 \quad (5)$$

Where Π is the canonical orthographic projection matrix, E a known external camera calibration matrix, and s the estimated per-frame scale. Although, given R this problem is

convex in a and s together³, for an unknown rotation matrix R the problem is extremely non-convex – even if a is known – and prone to sticking in local minima using first or second order gradient descent. The local optima often lie far apart in pose space and getting stuck in a poor optima leads to a significantly worse 3D reconstructions.

³To see this consider the trivial reparameterization where we solve for $s\mu + b \cdot \mathbf{e}$ and then let $a = b/s$.

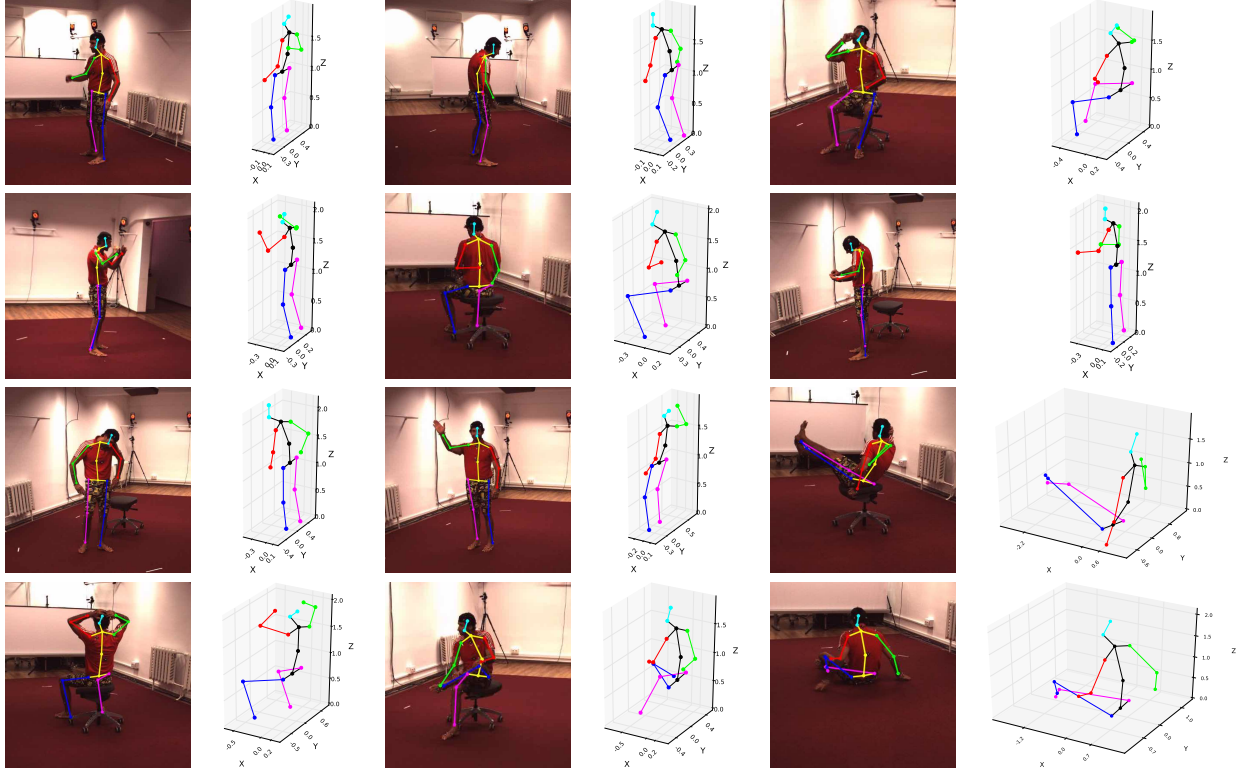


Figure 6: Results from the Human3.6M dataset. The identified 2D landmark positions and 3D skeleton is shown for each pose taken from different actions: Walking, Phoning, Greeting, Discussion, Sitting Down.

We take advantage of the matrix R 's restricted form that allows it to be parameterized in terms of a single angle θ (see (2)). Rather than attempting to solve this optimization problem using local gradient based methods we quantize over the space of possible rotations, and for each choice of rotation, we hold this fixed and solve for s and a , before picking the minimum cost solution of any choice of R . With fixed choices of rotation the terms $\Pi ER\mu$ and $\Pi ER\epsilon$ can be computed in advance, and finding the optimal a becomes a simple linear least square problem.

This process is highly efficient and by oversampling the rotations and exhaustively checking in 10,000 locations we can guarantee that a solution extremely close to the global optima is found. In practice, using 20 samples and refining the rotations and basis coefficients of the best found solution using a non-linear least squares solver obtains the same reconstruction, and we make use of the faster option of checking 80 locations and using the best found solution as our 3D estimate. This puts us close to the global optima and has the same average accuracy as finding the global optima. Moreover, it allows us to upgrade from sparse landmark locations to 3D using a single Gaussian at around 3,000 frames a second using python code on a standard laptop.

To handle models consisting of a mixture of Gaussians, we follow [26] and simply solve for each Gaussian independently and select the most probably solution.

5.4. Projecting 3D poses onto 2D belief maps

Our *projected pose model* is interleaved throughout the architecture (see Figure 1). The goal is to correct the beliefs regarding landmark locations at each stage, by fusing extra information about 3D physical plausibility. Given the solution R , s , and a from the previous component, we estimate a physically plausible projected 3D pose as

$$\hat{Y} = s\Pi ER(\mu + a \cdot \mathbf{e}) \quad (6)$$

which is then embedded in a belief map as

$$\hat{b}_{i,j}^p = \begin{cases} 1 & \text{if } (i,j) = \hat{Y}_p \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

and then convolved using Gaussian filters.

5.5. 2D Fusion of belief maps

The 2D belief maps predicted by our *probabilistic 3D pose model* are fused with the CNN-based belief maps b^p according to the following equation

$$f_t^p = w_t * b_t^p + (1 - w_t) * \hat{b}_t^p \quad (8)$$

where $w_t \in [0, 1]$ is a weight trained as part of the end-to-end learning. This set of fused belief maps f_t is then passed to the next stage and used as an input to guide the 2D re-estimation of joint locations, instead of the belief maps b_t used by convolutional pose machines.

	Directions	Discussion	Eating	Greeting	Phoning	Photo	Posing	Purchases
LinKDE [15]	132.71	183.55	132.37	164.39	162.12	205.94	150.61	171.31
Li <i>et al.</i> [21]	-	136.88	96.94	124.74	-	168.68	-	-
Tekin <i>et al.</i> [36]	102.39	158.52	87.95	126.83	118.37	185.02	114.69	107.61
Tekin <i>et al.</i> [34]	-	129.06	91.43	121.68	-	162.17	-	-
Tekin <i>et al.</i> [35]	85.03	108.79	84.38	98.94	119.39	95.65	98.49	93.77
Zhou <i>et al.</i> [48]	87.36	109.31	87.05	103.16	116.18	143.32	106.88	99.78
Sanzari <i>et al.</i> [30]	48.82	56.31	95.98	84.78	96.47	105.58	66.30	107.41
Ours - Single PPCA Model	68.55	78.27	77.22	89.05	91.63	110.05	74.92	83.71
Ours - Mixture PPCA Model	64.98	73.47	76.82	86.43	86.28	110.67	68.93	74.79
	Sitting	Sitting Down	Smoking	Waiting	Walk Dog	Walking	Walk Together	Average
LinKDE [15]	151.57	243.03	162.14	170.69	177.13	96.60	127.88	162.14
Li <i>et al.</i> [21]	-	-	-	-	132.17	69.97	-	-
Tekin <i>et al.</i> [36]	136.15	205.65	118.21	146.66	128.11	65.86	77.21	125.28
Tekin <i>et al.</i> [34]	-	-	-	-	130.53	65.75	-	-
Tekin <i>et al.</i> [35]	73.76	170.4	85.08	116.91	113.72	62.08	94.83	100.08
Zhou <i>et al.</i> [48]	124.52	199.23	107.42	118.09	114.23	79.39	97.70	113.01
Sanzari <i>et al.</i> [30]	116.89	129.63	97.84	65.94	130.46	92.58	102.21	93.15
Ours - Single PPCA Model	115.94	185.72	88.25	88.73	92.37	76.48	77.95	92.96
Ours - Mixture PPCA Model	110.19	173.91	84.95	85.78	86.26	71.36	73.14	88.39

Table 1: A comparison of our approach on the Human3.6M dataset against competitors that follow *Protocol 1* for evaluation. We substantially outperform all other methods in terms of average error showing a 4.7mm average improvement over our closest competitor. Note that some approaches [36, 48] use video as input instead of a single frame.

	Average error
Yasin <i>et al.</i> [44]	108.3
Rogez <i>et al.</i> [29]	88.1
Ours - Mixture PPCA Model	70.7

Table 2: A comparison of our approach on the Human3.6M dataset against competitors that follow *Protocol 2* for evaluation.

	Average error
Bogo <i>et al.</i> [7]	82.3
Ours - Mixture PPCA Model	79.6

Table 3: A comparison of our approach on the Human3.6M dataset against Bogo *et al.* [7] who follow *Protocol 3* described in Section 6. The comparison is done using the same set of joints as Bogo *et al.* [7] (14 rather than 17).

Final lifting: The belief maps produced as the output of the final stage ($t = 6$) are then lifted into 3D to give the final estimate for the pose (see figure 1) using our algorithm to lift 2D poses into 3D.

6. Experimental evaluation

Human3.6M dataset: Our model was trained and tested on the Human3.6M dataset consisting of 3.6 million ac-

curate 3D human poses [15]. This is a video and mocap dataset of 5 female and 6 male subjects, captured from 4 different viewpoints, that show them performing typical activities (talking on the phone, walking, greeting, eating, etc.).

2D Evaluation: Figure 5 shows how the 2D predictions are improved by our projected pose model, reducing the overall mean error per landmark. The 2D error reduction using our full approach over the estimates of [43] is comparable in magnitude to the improvement due to the change of architecture moving from the work Zhou *et al.* [48] to the state-of-the-art 2d architecture [43] (i.e. a reduction of 0.59 pixels vs. 0.81 pixels). See Table 4 for details.

3D Evaluation: Several evaluation protocols have been followed by different authors to measure the performance of their 3D pose estimation methods on the Human3.6M dataset. Tables 1, 2 and 3 show comparisons of the performance of our 3D pose estimation with previously published works, where we take care to evaluate using the appropriate protocol in each case.

Protocol 1, the most standard evaluation protocol on Human3.6M was followed by [15, 21, 36, 34, 35, 48, 30]. The training set consists of 5 subjects (S1, S5, S6, S7, S8), while the test set includes 2 subjects (S9, S11). The original frame rate of 50 fps is downsampled to 10 fps and the evaluation is on sequences coming from all 4 cameras and all trials. The reported error metric is the *3D error* which corresponds to the average Euclidean distance of the estimated 3D joints to

the ground truth. The error is averaged over all 17 joints of the Human3.6M skeletal model. Table 1 shows a comparison between our approach and other competing approaches that use *Protocol 1* for evaluation.

Our baseline method using a single unimodal probabilistic PCA model outperforms almost every method in most action types, with the exception of Sanzari *et al.* [30], which it still outperforms on average across the entire dataset. Our mixture model improves on this again, offering a 4.76mm improvement over Sanzari *et al.*, our closest competitor.

Protocol 2, followed by [44, 29], selects 6 subjects (S1, S5, S6, S7, S8 and S9) for training and subject S11 for testing. The original video is downsampled to every 64th frame and evaluation is performed on sequences from all 4 cameras and all trials. The error metric reported in this case is the *3D pose error* equivalent to the per-joint 3D error up to a similarity transformation (i.e. each estimated 3D pose is aligned with the ground truth pose, on a per-frame basis, using Procrustes analysis). The error is averaged over a subset of 14 joints. Table 2 shows a comparison between our approach and other competing approaches that use *Protocol 2* for evaluation.

Although in this specific case, our model had been trained using only the 5 subjects used for training in *Protocol 1* (one fewer subject), our model still outperforms both other methods [44, 29].

Protocol 3, followed by [7] selects the same subjects for training and testing as *Protocol 1*. However, evaluation is only on sequences captured from the frontal camera (cam 3) from trial 1 and the original video is not sub-sampled. The error metric used in this case is the *3D pose error* as described in Protocol 2. The error is averaged over 14 joints. Table 3 shows a comparison between our approach and [7].

Our method outperforms Bogo *et al.* [7] by almost 3mm on average. It is important to remind the reader that, unlike our approach, Bogo *et al.* [7] exploits a high-quality detailed statistical 3D body model [22] trained on thousands of 3D body scans, that captures both the variation of human body shape and its deformation through pose.

MPII and LEEDS dataset: Our proposed approach trained exclusively on the Human3.6M dataset can be used to identify 2D and 3D landmarks of images contained in different datasets. Figure 7 shows some qualitative results on the MPII dataset [5], whereas Figure 8 shows qualitative results on the LEEDS dataset [17], including failure cases. Notice how our *probabilistic 3D pose model* generates anatomically plausible poses even though the 2D landmark estimations are not all correct. However, as shown in bottom row, even small errors in 2D pose can lead to drastically different 3D poses. These inaccuracies could be mitigated without further 3D data by annotating additional RGB images for training from different datasets.

	2D pixel error
Zhou <i>et al.</i> [48]	10.85
Trained CPM [43] architecture	10.04
Ours including 3D refinement	9.47

Table 4: 2D errors of related work. Our approach of lifting the 2D landmark predictions into a plausible 3D model and then projected back into image substantially reduces the error. Note that [48] use video as input and knowledge of the action label.

7. Conclusion


We have presented a novel approach to human 3D pose estimation from a single image that outperforms previous solutions. To our knowledge we are the first to approach this as a problem of iterative refinement in which 3D proposals help refine and improve upon the 2D estimates. Our approach shows the importance of thinking in 3D even for 2D pose estimation within a single image, with our iterative 3D model demonstrating better 2D accuracy than Convolutional Pose Machines [43], the iterative 2D approach it is based upon.

Our novel approach for upgrading from 2D to 3D is extremely efficient, and runs in CPU-based python at around 1,000 frames a second, while a GPU-based real-time approach for Convolutional Pose Machines has been announced. Integrating these two systems to provide a reliable real-time 3D pose estimator seems like a natural future direction, as does integrating our approach with a simpler 2D approach for real-time pose estimation on lower power devices.

References

- [1] A. Agarwal and B. Triggs. Recovering 3d human pose from monocular images. *IEEE transactions on pattern analysis and machine intelligence*, 28(1):44–58, 2006. 2
- [2] S. Agarwal, K. Mierle, and Others. Ceres solver. <http://ceres-solver.org>. 4
- [3] I. Akhter and M. J. Black. Pose-conditioned joint angle limits for 3d human pose reconstruction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1446–1455, 2015. 2
- [4] I. Akhter, Y. Sheikh, S. Khan, and T. Kanade. Trajectory space: A dual representation for nonrigid structure from motion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2011. 2
- [5] M. Andriluka, L. Pishchulin, P. Gehler, and B. Schiele. Human pose estimation: New benchmark and state of the art analysis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014. 9, 10
- [6] C. Barrón and I. A. Kakadiaris. Estimating anthropometry and pose from a single uncalibrated image. *Computer Vision and Image Understanding*, 81(3):269–284, 2001. 2
- [7] F. Bogo, A. Kanazawa, C. Lassner, P. Gehler, J. Romero, and M. J. Black. Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In *European*

Success cases



Failure cases





Figure 7: Results on images from the MPII [5]. Our model was not trained on images as diverse as those contained in this dataset, however it often retrieves correct 2D and 3D joint positions. The last row shows example cases where the method fails either in the identification of 2D or 3D landmarks.

- Conference on Computer Vision*, pages 561–578. Springer, 2016. 3, 8, 9
- [8] C. Bregler, A. Hertzmann, and H. Biermann. Recovering non-rigid 3d shape from image streams. In *Computer Vision and Pattern Recognition, 2000. Proceedings. IEEE Conference on*, volume 2, pages 690–696. IEEE, 2000. 2, 4
- [9] X. Chen and A. L. Yuille. Articulated pose estimation by a graphical model with image dependent pairwise relations. In *Advances in Neural Information Processing Systems*, pages 1736–1744, 2014. 3
- [10] J. Cho, M. Lee, and S. Oh. Complex non-rigid 3d shape recovery using a procrustean normal distribution mixture model. *International Journal of Computer Vision*, 117(3):226–246, 2016. 2
- [11] C. H. Ek, P. H. S. Torr, and N. D. Lawrence. Gaussian process latent variable models for human pose estimation. In A. Popescu-Belis, S. Renals, and H. Bourlard, editors, *MLMI*, volume 4892 of *Lecture Notes in Computer Science*, pages 132–143. Springer, 2007. 2
- [12] A. Elgammal and C. Lee. Inferring 3d body pose from silhouettes using activity manifold learning. In *CVPR*, 2004. 2
- [13] X. Fan, K. Zheng, Y. Zhou, and S. Wang. Pose locality constrained representation for 3d human pose reconstruction. In *European Conference on Computer Vision*, pages 174–188. Springer, 2014. 2
- [14] P. Gotardo and A. Martinez. Computing smooth time trajectories for camera and deformable shape in structure from motion with occlusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2011. 2
- [15] C. Ionescu, D. Papava, V. Olaru, and C. Sminchisescu. Human3. 6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE transactions on pattern analysis and machine intelligence*, 36(7):1325–1339, 2014. 4, 8
- [16] A. Jain, J. Tompson, M. Andriluka, G. W. Taylor, and C. Bregler. Learning human pose estimation features with convolutional networks. *arXiv preprint arXiv:1312.7302*, 2013. 3

Success cases



Failure cases




Figure 8: Results on images from the LEEDS [17] dataset. Our model was not trained on images as diverse as those contained in this dataset, however it often retrieves correct 2D and 3D joint positions. The last row shows example cases where the method fails either in the identification of 2D or 3D landmarks.

- [17] S. Johnson and M. Everingham. Clustered pose and nonlinear appearance models for human pose estimation. In *Proceedings of the British Machine Vision Conference*, 2010. doi:10.5244/C.24.12. 9, 11
- [18] H.-J. Lee and Z. Chen. Determination of 3d human body postures from a single view. *Computer Vision, Graphics, and Image Processing*, 30(2):148–168, 1985. 2
- [19] M. Lee, J. Cho, and S. Oh. Procrustean normal distribution for non-rigid structure from motion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2016. 2
- [20] S. Li and A. B. Chan. 3d human pose estimation from monocular images with deep convolutional neural network. In *Asian Conference on Computer Vision*, pages 332–347. Springer, 2014. 2, 3
- [21] S. Li, W. Zhang, and A. B. Chan. Maximum-margin structured learning with deep networks for 3d human pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2848–2856, 2015. 2, 3, 8
- [22] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black. Smpl: A skinned multi-person linear model. *ACM Transactions on Graphics (TOG)*, 34(6):248, 2015. 3, 9
- [23] G. Mori and J. Malik. Recovering 3d human body configurations using shape contexts. *PAMI*, 2006. 2
- [24] V. Parameswaran and R. Chellappa. View independent human body pose estimation from a single perspective image. In *Computer Vision and Pattern Recognition, 2004. CVPR 2004. Proceedings of the 2004 IEEE Computer Society Conference on*, volume 2, pages II–16. IEEE, 2004. 2
- [25] T. Pfister, J. Charles, and A. Zisserman. Flowing convnets for human pose estimation in videos. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1913–1921, 2015. 3
- [26] N. Pitelis, C. Russell, and L. Agapito. Learning a manifold as an atlas. In *Computer Vision and Pattern Recognition (CVPR)*, 2013 *IEEE Conference on*, pages 1642–1649. IEEE, 2013. 4, 7
- [27] V. Ramakrishna, T. Kanade, and Y. Sheikh. Reconstructing 3d human pose from 2d image landmarks. In *European Conference on Computer Vision*, pages 573–586. Springer, 2012. 2
- [28] V. Ramakrishna, D. Munoz, M. Hebert, J. A. Bagnell, and Y. Sheikh. Pose machines: Articulated pose estimation via inference machines. In *European Conference on Computer Vision*, pages 33–47. Springer, 2014. 5
- [29] G. Rogez and C. Schmid. Mocap-guided data augmentation for 3d pose estimation in the wild. In *Advances in Neural Information Processing Systems*, pages 3108–3116, 2016. 8, 9
- [30] M. Sanzari, V. Ntouskos, and F. Pirri. Bayesian image based 3d pose estimation. In *European Conference on Computer Vision*, pages 566–582. Springer, 2016. 3, 8, 9
- [31] L. Sigal, R. Memisevic, and D. Fleet. Shared kernel information embedding for discriminative inference. In *CVPR*, 2009. 2
- [32] E. Simo-Serra, A. Ramisa, G. Alenyà, C. Torras, and F. Moreno-Noguer. Single image 3d human pose estimation from noisy observations. In *Computer Vision and Pattern Recognition (CVPR)*, 2012 *IEEE Conference on*, pages 2673–2680. IEEE, 2012. 3
- [33] C. J. Taylor. Reconstruction of articulated objects from point correspondences in a single uncalibrated image. In *Computer Vision and Pattern Recognition, 2000. Proceedings. IEEE Conference on*, volume 1, pages 677–684. IEEE, 2000. 2
- [34] B. Tekin, I. Katircioglu, M. Salzmann, V. Lepetit, and P. Fua. Structured prediction of 3d human pose with deep neural networks. In *British Machine Vision Conference (BMVC)*, 2016. 3, 8

- [35] B. Tekin, P. Márquez-Neila, M. Salzmann, and P. Fua. Fusing 2d uncertainty and 3d cues for monocular body pose estimation. *arXiv preprint arXiv:1611.05708*, 2016. 8
- [36] B. Tekin, X. Sun, X. Wang, V. Lepetit, and P. Fua. Predicting people’s 3d poses from short sequences. *arXiv preprint arXiv:1504.08200*, 2015. 8
- [37] M. E. Tipping and C. M. Bishop. Probabilistic principal component analysis. *Journal of the Royal Statistical Society, Series B*, 61:611–622, 1999. 4, 5
- [38] C. Tomasi and T. Kanade. Shape and motion from image streams under orthography: a factorization method. *International Journal of Computer Vision*, 9(2):137–154, 1992. 4
- [39] J. J. Tompson, A. Jain, Y. LeCun, and C. Bregler. Joint training of a convolutional network and a graphical model for human pose estimation. In *Advances in neural information processing systems*, pages 1799–1807, 2014. 2, 3
- [40] A. Toshev and C. Szegedy. Deeppose: Human pose estimation via deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1653–1660, 2014. 2, 3
- [41] C. Wang, J. Flynn, Y. Wang, and A. L. Yuille. Representing data by a mixture of activated simplices. *arXiv preprint arXiv:1412.4102*, 2014. 4
- [42] C. Wang, Y. Wang, Z. Lin, , A. Yuille, and W. Gao. Robust estimation of human poses from a single image. In *Computer Vision and Pattern Recognition (CVPR)*, 2014. 2
- [43] S.-E. Wei, V. Ramakrishna, T. Kanade, and Y. Sheikh. Convolutional pose machines. *arXiv preprint arXiv:1602.00134*, 2016. 1, 3, 5, 8, 9
- [44] H. Yasin, U. Iqbal, B. Kruger, A. Weber, and J. Gall. A dual-source approach for 3d pose estimation from a single image. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 8, 9
- [45] R. Zhao, Y. Wang, and A. Martinez. A simple, fast and highly-accurate algorithm to recover 3d shape from 2d landmarks on a single image. *arXiv preprint arXiv:1609.09058*, 2016. 2
- [46] X. Zhou, X. Sun, W. Zhang, S. Liang, and Y. Wei. Deep kinematic pose regression. *arXiv preprint arXiv:1609.05317*, 2016. 3
- [47] X. Zhou, M. Zhu, S. Leonardos, and K. Daniilidis. Sparse representation for 3d shape estimation: A convex relaxation approach. *arXiv preprint arXiv:1509.04309*, 2015. 2
- [48] X. Zhou, M. Zhu, S. Leonardos, K. Derpanis, and K. Daniilidis. Sparseness meets deepness: 3d human pose estimation from monocular video. *arXiv preprint arXiv:1511.09439*, 2015. 3, 5, 8, 9